ʼʼʼʼʼATION PAGE

AD-A277 424

1a. REPORT SECURITY CLAS Unclassified	RICTIVE MARKINGS
2a. SECURITY CLASSIFICAT	RIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	Unlimited

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR· 94 0126

6a. NAME OF PERFORMING ORGANIZATION Department of Psychology	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research/NL

6c. ADDRESS (City, State and ZIP Code) University of Illinois 506 S. Wright Street Urbana, IL 61801	7b. ADDRESS (City, State and ZIP Code) Bolling AFB, DC 20332-6448

8a. NAME OF FUNDING/SPONSORING ORGANIZATION same as 7a	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR- 89- C447

8c. ADDRESS (City, State and ZIP Code) same as 7b	10 SOURCE OF FUNDING NOS.			
	PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.
11. TITLE (Include Security Classification) Reminding-based Learning	61102F	2313	A4	

12. PERSONAL AUTHOR(S)
Brian H. Ross

13a. TYPE OF REPORT Annual Technical	13b. TIME COVERED FROM 1/21/93 TO 1/20/94	14. DATE OF REPORT (Yr., Mo., Day) February 8, 1994	15. PAGE COUNT 11

16. SUPPLEMENTARY NOTATION

17.	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB. GR.	
05	09		

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

When learning new cognitive skills involving problem solving, novices are often reminded of earlier problems. The use of earlier problems is a common means of problem solving and affects the learning of the skill. This project has three aims in understanding this learning. First, the representation of the resulting generalizations is being examined. Generalizations formed from remindings are likely to be conservative, in that they may be more tied to the examples than many current theories allow. A main aim of the project is to distinguish and test different forms of this conservatism. Second, the development of problem solving expertise is examined by focusing on differences in how typical and atypical problems are solved. Third, the effects of such reminding-based learning in everyday problem solving is examined to extend the findings and test some theoretical ideas that are difficult to investigate in more formal domains. This report provides an overview of this work and the progress on these objectives during the last year.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS ☐	21. ABSTRACT SECURITY CLASSIFICATION Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Tangney	22b. TELEPHONE NUMBER (Include Area Code) (202)767-5021	22c. OFFICE SYMBOL NL

DD FORM 1473, 83 APR EDITION OF 1 JAN 73 IS OBSOLETE. **UNCLASSIFIED**

SECURITY CLASSIFICATION OF TH··

ANNUAL TECHNICAL REPORT: · REMINDING-BASED LEARNING

In this annual report, I have again repeated the overview and background needed to understand the progress made since the last report. Although this strategy makes this report partially redundant with the 1993 report, it allows the reader to understand this report without referring back to the earlier report for clarification. I have only included projects that have been worked on in the last year.

OBJECTIVES

When learning a new cognitive skill, novices spend much of their time solving problems. In doing so, it is common for novices to think back to an earlier problem that the current problem reminds them of and use this earlier problem to help solve the current problem (e.g., Ross, 1984, 1987, 1989a). This use of the earlier problem not only affects performance on the current problem, but also provides the learner with additional knowledge that can be accessed and used on later problems. The aim of this research is to understand the nature of the learning that results from this use of earlier problems. Little is known about this crucial source of learning.

Such within-domain analogies occur frequently during learning. In the view presented in this work, a generalization is formed from making an analogy between problems. Rather than positing a separate generalization process that operates upon completed instances, the generalization may be a byproduct of the analogy. In using the earlier problem to help solve the current problem, comparisons must be made and some aspects generalized over. Remindings, by setting up the analogy, may determine what pairs of problems are compared and, hence, what generalizations are made. The learning comes about because, while the noticing might be based on a variety of similar aspects (including superficial ones) between the problems, the comparisons forces the generalization of many of the aspects. My earlier work (e.g., Ross & Kennedy, 1990) has shown that the use of earlier problems allows novices to begin to form generalizations across problems. Thus, this means of learning may be one way in which novices can begin to develop more expert-like knowledge structures. However, much work remains to be done to understand this learning. The research I have been conducting has three goals.

The first goal is to understand the nature of the resulting generalization, the information included in its representation. If the generalization results from the reminding, it is likely to be a conservative one, in that it will be somewhat tied to the problems from which it arises. In fact, most theories of learning assume some conservatism, but my research (a) distinguishes among different types of assumptions of conservatism, (b) relates these distinctions to current theories, and (c) tests these distinctions. Thus, this work examines the specificity of what is learned, as well as its generality.

The second goal of this research is to examine the implications of these ideas for the development of expertise. Research on expertise in mathematics and physical science domains suggests that experts have problem schemas that allow them to categorize problems, as well as associated procedures for solving problems of that type (e.g., Chi, Feltovich, & Glaser, 1981). Despite the importance of these schemas, little research examines how they are learned. A common idea is that they may develop from the comparison of problems. However, this idea leads to two questions? One, how do people know which problems to compare? Two, why are people comparing problems (i.e., what is the nature of this comparison)? The reminding-based learning view suggests that people compare problems when one makes them think back to another and they do so in order to use the earlier one to solve the current problem. In addition, this view suggests that people may develop problem schemas that are influenced by the superficial aspects, because these aspects are known to affect remindings. Thus, part of this project is concerned with the development of such problem schemas and the possibility that some schemas may critically depend upon superficial contents, even in experts.

94-09341

94 3 25 054

The third goal of this project is to begin to extend this work to more everyday problem solving situations. This extension will not only allow the application and test of these ideas in an important new setting, it will also force the extension of this work to important situations that are hard to experiment with in more formal domains.

STATUS OF PROJECTS

In this section, I will provide some details of the status of the projects underway. For each goal, I will first give a brief summary of the findings and then provide further details.

1. Conservatism of learning

As described earlier, this project examines the representation of the resulting generalization. The focus here has been on asking how the problem solving affects what gets incorporated into the generalization and the effect on later performance.

Summary. One project has been examining how the representation of the generalization may depend upon the details of how the earlier problem is used. The early work in this project, reported last year, has not only led to further research along the same lines, but has also spawned a rather different approach to problem categorization. A second project investigates the interactions among the different processes involved in reminding, analogical transfer, and generalization.

1a. The effect of problem use on the generalization. In addition to testing the conservatism of the generalization in the way described above, a central idea of the reminding-based view is that the generalization depends upon the earlier problem comparison. This selectivity effect has been investigated by me in the context of category learning (Ross, Perkins, & Tenpenny, 1990), because in formal domains different earlier problems will often lead to very similar generalizations. However, in this study I examine whether the details of how the problem is used (rather than which problem) affects what is learned. As will be seen, it can be viewed as a type of transfer-appropriate processing in problem solving. In (i), I report on some earlier work and the extensions in the last year. In (ii), I report on a new project that approaches this general idea from a different perspective and the initial results of some experiments.

(i). *Problem categorization:* The basic paradigm for examining the effect of using an earlier problem is referred to as the cuing method. The typical experiment involves study examples, first test problems (with cues to the study examples), and second test problems. This simple case allows us to isolate the specific influences on each test and to detail the learning on a step-by-step level. More specifically, subjects are instructed in a simplified formal domain, elementary probability theory. They are given a short introduction to some probability concepts and then learn four principles (e.g., permutations). For each principle, an explanation of the principle with the appropriate formula is given. The subject then helps to solve a word problem that requires the use of the principle. This study problem is presented in a workbook format, with the subjects guided through the solution. After learning these principles (4 mins/principle), subjects are given first test problems to solve (3 mins). In many of the studies, some or all of these first test problems include a cue (e.g., "this is like the earlier golf problem"), which has been shown to result in a generalization (Ross & Kennedy, 1990). After each problem, feedback is given and subjects study the solution. A second test is then given. Of main interest is performance on the second test problem as a function of the experimental manipulation. In all of these experiments, principles are rotated through conditions and several test examples are used for each principle to avoid effects due to any particular example.

The earliest studies using the cuing paradigm showed that cuing on the first test led to

improved performance on the second test in three ways: better problem solving performance, higher probability of accessing the appropriate formula, and better instantiation of a provided formula (Ross & Kennedy, 1990, Exs. 1 and 2). After conducting these studies, I also became interested in problem categorization, which is thought to be a crucial element in developing expertise. Thus, I used the cuing paradigm but instead of a second test problem, subjects were given 12 problems (3 of each principle) and asked to categorize them by principle. That is, at the first test, half of the principles were cued and half were not. If cuing affects problem categorization, then higher categorization should occur for the cued principles. In one study, I gave them formulae to use as the categories for sorting and in another I gave them the principle names. The (unexpected) result was that the cuing had no effect.

In an earlier annual report, I reported a study that investigated one difference between the cue method and the usual way in which people think of earlier problems. Usually when people are reminded they need to decide whether the reminding is appropriate. That is, they need to decide whether to use the reminding or not, since we often get reminded of earlier problems for inappropriate reasons. This difference may be crucial because a determination of appropriateness may include deciding whether the problems are of the same "type" so that one could be used to help solve another. Thus, such a determination might affect what is learned about problem categories.

The study used the same procedure and materials that led to no cuing effect and made one change. On the principles that subjects are cued, they are told there is a chance that these cues are not correct. Their task is to first decide if the cue is correct or not. They are then given feedback on this response and given the correct cue. In fact, all the cues are correct so they are given exactly the same cues as in the earlier study, but they have to determine appropriateness before using the cued problem. The results in this case showed an effect of cuing, with the cued principles leading to .47 of the second test problems being correctly categorized compared to .30 for the not cued principles, t(29) = 2.17.

In my last report, I mentioned a replication of this which also included the usual cues. The results were encouraging: the probabilistic cue subjects showed a small, but consistent and significant effect of cuing, .47 vs. .40, t(47)=2.77. The subjects receiving the usual cues, showed a slight, nonsignificant, effect in the other direction, .413 vs. .455.

In the last year, my graduate RA, Matthew Kilbane, and I have replicated this result for the probabilistic cue subjects (.48 vs. .43). Our next experiment, which is in progress, attempts to provide a deeeper understanding of the reason for the problem similarity by providing a further description of each test problem phrased in a way to highlight the relevant information. The point of this manipulation is to ask whether the cuing to an earlier problem allows people to make even better use of this relevant information. We have two motivations for this manipulation change. First, we wanted to show that this probabilistic cue effect can occur even when the uncued conditions lead to higher performance. Second, and most importantly, we think this manipulation will have important implications for how complex problem categories might best be instructed. That is, while the cuing helped performance in our earlier experiment, the effect on problem categorization was small. Our analysis suggested that what was needed along with the comparison was some way of pointing learners to the reason behind the problem similarity, which is what we are trying to do with the current cues. Assuming it works, we will have good evidence that how the earlier problem is used affects what is learned.

(ii) *Problem categorization as a function of problem use.* One new project begun this year came out of the ideas for the problem categorization study. First, the work in problem solving suggests that problem categorization is among the most crucial aspects of problem solving and one of the most difficult to learn. (One way to think about this is in terms of the applicability conditions used in production systems, which most computational modelers have argued is the

toughest part of getting the systems to work). Second, the work in categorization focuses on the instances being categorized as if the goal was to know the category rather than use the category to do something with the instance (e.g., make a prediction, inference, etc.). In this project, I take the problem solving perspective and try to combine it with the project mentioned in (i) to show how the instances are used affects what is learned about the category. My preliminary work on this uses simple algebra equations and contrasts what is learned about the categories in terms of the effect on later categorization of more complex instances. I will mention two results showing this type of effect to make clearer the idea and directions I hope to take it.

(iia) Contrasting categorization and use. I used two simple types of algebra equations that differ in terms of the order of the operations used to solve them. For example, $a + (bx/c) = p$ can be solved for x by subtracting a, mulitplying by c, and dividing by b, which I shall call SMD (for this experiment it does not matter if they go SDM instead). A second equation type, for example, $(q + mx)/b = s$, can be solved by multiply, subtracting, then dividing (MSD). In addition, I confounded some mathematically irrelevant properties with each type, such as the letters likely to be used as constants (early in the alphabet vs. late) and the parentheses (whether it excluded part of the sum or not). Two groups learned to categorize these equations as type 1 or type 2. One group simply categorized them and were given feedback. The second group also had to solve the equations for x. The results show that both groups are excellent at categorizing new instances like the study ones, but the question of interest is what happens when the test items separate out the effects of operations, letters, and parentheses. The simple story is that the solving group is much more likely to choose on the basis of operations, while the categorization alone group is more greatly drawn by the mathematically irrelevant properties. Although this may not seem like a surprising result, the point to remember is that (1) these were very simple equations for which all subjects had much experience and the two types had obvious differences and (2) it demonstrates the importance of the use in learning the problem categories.

(iib) Different uses lead to different generalizations. One difficulty with this first experiment is that it is hard to rule out a simple better learning explanation for the performance of the solving group. So, in the second experiment, two groups solved the problems differently and ended up with corresponding differences in their understanding of the problem categories. In particular, I added in a y variable, such as in the equation $9 + sx/t = (5qy + b)/f$. One group solved for x (making it SMD), while the other group solved for y (making it MSD). The other types of equations were MSD for the x group and SMD for the y group. Thus, though they saw the same equations and categorized them in the same way at study (i.e., the equation given was called Type 1 by both groups), they are predicted to be learning very different things about what characterizes each type. This prediction was confirmed by 85% of the test items (in which the variable z was used).

These experiments provide an important direction on my work on problem categorization for two reasons. First, they provide a way of examining more directly the selectivity issue, how the use of the problems affect what is learned from them, an idea of importance for both theoretical and instructional reasons. Second, they provide an important bridge between the huge amount of work on categorization to the much smaller body of work on problem categorization.

1b. Interactions among different processes. The work on selectivity can be viewed in a slightly different way - the generalization depends upon the details of the problem solving. This idea has led me to go back and consider in much more detail how the different processes involved in reminding, application, and generalization may be interacting. As I mentioned in last year's report, the first studies I conducted in this project are in collaboration with Prof. Gary Bradshaw, a faculty member at University of Colorado who is now permanently at Illinois We have focused on remindings and transfer. All current views of analogical problem solving (including my own) assume the new (target) problem is represented and that this representation leads to a retrieval of

some earlier problem (which is then mapped and transferred to the new target problem). However, there is also evidence that people may be reminded during the reading of a problem.

Our studies examine whether such remindings of earlier problems may mean that the target problem's representation is affected by the reminding. The importance of this point is that much of the focus on mapping as a separate process may be misplaced, IF the earlier problem is affecting how the target problem is represented (i.e., much of the transfer occurs DURING the target being represented, not between the represented target and the earlier problem). So, if this problem makes me think of an earlier one, I may fit the current problem to the earlier one. Bradshaw and I have begun to test this idea in a general way, using how simple stories are interpreted. Our results have been encouraging: we find a simple cue (a proper name) reminds people of an earlier story and affects their interpretation of the current one. In addition, we find that this is an encoding effect in that one can see the effect on sentence-by-sentence reading times.

Last year, I reported that we had just written this work up and submitted it as a paper to the Cognitive Science Conference. Not only was it accepted there, but a more extensive write-up has been accepted for publication in *Memory & Cognition* (unusually, with suggestions for only minor changes, which have been done). After these general interpretation effects, we will now begin to look at how the particulars of the target representation may be affected by the reminding. I am quite excited by this project and think it will be an important addition to our understanding of the reminding effects on problem solving and learning.

2. The development of problem solving expertise

As discussed earlier, reminding-based learning provides one perspective on the nature of problem schemas and their development. The work already discussed can be viewed as very early precursors to problem schemas, but no evidence has yet been presented that in fact they do lead to problem schemas.

My focus has been on the specialized schemas that often include superficial information as well. To study this, I have written a number of algebra word problems (e.g., distance, interest, mixture). Each problem has two versions: one with appropriate (typical) contents and one with neutral contents. To make clear the advances in this project, I need to review the earlier work I reported in last year's progress report. The first two studies conducted have followed up observations noted in a chapter a number of years ago (Hinsley, Hayes, & Simon, 1977). First, we have protocols of subjects solving these problems in order to ask how they might differentially solve the appropriate and neutral problems. Hinsley et al. argued that the appropriate ones were solved by schema instantiation, while the others were solved by translating each sentence to an equation. This result is quite important, but their results were based on few observations and no quantitative data were provided. However, we were unable to replicate this finding, not because it is incorrect, but because so few of the problems were solved at all (less than 50%), making it impossible to see differences between conditions. We considered several training studies, but opted for what I think will be a more fruitful tack. An honors student, Steven Blessing (now at Carnegie-Mellon, working with John Anderson), has redone this experiment with three major changes. First, we have used high math students (UI students who graduated from a math and science academy). Second, he has included a set of inappropriate problems as well, that is problems of a given type with the story contents of a different type. Third, we have examined the time and proportion correct. Our finding is that the appropriate and neutral conditions do not differ, but they are both more accurate and faster than the inappropriate condition. In addition, he has collected problem solving protocols from another group of subjects.

In the second study originally conducted, other subjects received the same problems one clause at a time and were required after each clause to say what type of problem it is. Hinsley et al.

showed people could do this quite readily. Our interest is in asking how this measure may differ for appropriate and neutral contents. The appropriate problems are categorized much earlier.

We conducted several other studies. First, the high math subjects were tested in the clause study with appropriate, neutral, and inappropriate problems. We found that the appropriate were categorized faster than the neutral (as before) and the inappropriate were categorized slower than the neutral. The inappropriate contents often lead to subjects initially making the (content-appropriate) inappropriate categorization, before "seeing" the right category. Second, the regular subjects (i.e., UI students not from the math academy) participated in the timed problem solving task. We found the same results as with the academy students (though the means were lower). Together, these studies provide important information about the problem schemas and their use, particularly concerning the inclusion and use of superficial aspects. Our current hypothesis is that the usual or appropriate contents allow one to more quickly categorize problems (i.e., the clause result), but that even neutral problems are generally categorized correctly by the end of the problem (another result from the clause study). Thus, there is little accuracy or latency difference in problem solving between these two conditions. However, the inappropriate problems are more difficult to categorize and the categorization is often uncertain even after the problem has been read.

However, we had not seen any evidence that the appropriate materials provide any benefit to how the problem is solved. Because our hypothesis is that for these simple problems, the schema is recognized by the end of the reading even for the neutrals, we wanted some manipulation in which the schema might provide some benefit if it were recognized early. We decided to use more complex problems, in which we added irrelevant information. The idea is that if the solver has the appropriate schema, s/he may be better able to decide on the relevance of the different information in the problem. This manipulation was quite successful. First, a group of the academy subjects were faster and more accurate in solving the appropriate problems versus the neutral problems (the first time this difference has come out). Second, Blessing collected some protocols and found that the appropriate problems led to subjects talking more about relevant information compared to the neutral problems (i.e., a higher proportion of the protocol statements concerned information that was relevant to the problem solution).

We were encouraged by these findings and believe that they may allow us to see some benefit of the content on how the problem is solved (not just on how the schema may be accessed). Last year, I mentioned that a likely possibility for the next step was to get protocols from some real experts (algebra teachers) and examine them for whether the contents are used throughout the problem solving. That is what I did, testing the experts with the same complex problems used earlier. The results were quite interesting. First, showing us they were experts, all the problems were solved correctly. Second, when the protocols were examined for the sensitivity to the relevance of the information, we found exactly the same pattern as with the experienced subjects. Third, the protocols suggested (though only suggested) that the solution of the problem was affected by the content in its details. For example, one subject solving a neutral problem (about a boy going up and down a down escalator) actually ended up talking about it as if it had the contents typically associated with problems of this type (a river current problem in which a boat goes down river with the current and then upriver against the current).

As a whole, these experiments provide the clearerst evidence for the effect of problem content on the performance of experienced problem solvers. We have just submitted a paper based on these results to the Cognitive Science Conference and are about to submit a more extensive version for publication.

3. Everyday problem solving

In addition to the work on probability theory and formal domains, I have been examining how remindings may be used in less formal situations. In particular, I have been examining how

they may affect the simple categories that are learned. I view this work as parallel to the problem solving work, but it sometimes allows me to more cleanly investigate certain issues, especially selectivity effects (Ross, Perkins, & Tenpenny, 1990). In addition, I have been taking this work further (and relating it to the category use ideas mentioned in 1aii) to examine how categories are used in making predictions.

Summary. Two main projects have been conducted under this idea. First, the work on reminding-based generalizations in categorization has been completed and accepted for publication. I provided a greatly detailed review of this in last year's report and have sent in the preprint of this work, so I will only provide a short summary of the idea in 3a. The second project is also investigating the use of categories, with a focus on how categories may be used for prediction. Although this is a little more tangential to the reminding-based learning work, it blends in nicely with the effects of category use on learning (1aii) and also provides a more detailed assessment of how the categories are being used in a complex situation.

3a. Reminding based generalizations in categorization. The work from the last three years (with a graduate student RA, Tom Spalding) has extended my earlier work to show (1) that reminding-based generalizations occur in a common categorization study paradigm, (2) that remindings serve to **focus** the learner on to relevant features, (3) that these manipulations lead to differences in perceived **frequency**, (4) that they also affect people's **prediction** of the values for missing features, and (5) that the effects of early focus features affects what is learned from later instances. I believe that this idea may help to provide one means by which categories can be learned and help relate such learning to performance issues. In addition to a chapter based on these ideas (Ross & Spalding, 1991), Spalding and I have just had a paper on this accepted in *Journal of Experimental Psychology: Learning, Memory, and Cognition* (Spalding & Ross, in press) and have presented this work at MPA, Psychonomics, and APS. I gave a detailed report of this in last year's progress report and will not repeat it here because we have mainly worked on tying up loose ends in this project for publication in the last year. Again, the motivation is to bring the category work to bear on issues we have been examining in the problem solving domain. For example, problem categories (as discussed in section 2 of this report) are viewed as essential in problem solving, but it is not enough to categorize a problem. One must then use the associated procedures for solving the problem. Similarly, within the concept learning work, we are going beyond the question of how an instance is categorized to investigate how the categorization is used. These studies add not only to the idea of reminding-based learning, but also to how categories might be learned.

3b. Predictions from uncertain categorizations. My current work, along with a faculty member here, Gregory Murphy, has taken the ideas about the use of categories and examined it more directly in how the categories are used to make judgments. The next direction in this work is to extend it to simple problem solving situations. In many situations in which people need to make a prediction or judgment about an object, it is assumed that they access the category information about the object and use this, along with any specific knowledge they might have about the object, to make the prediction or judgment. However, often the categories are not certain, but rather people have only a likely category. An animal seen briefly is probably a cat, but it could be a dog or raccoon. How do people use this uncertainty in making predictions?

Our work on this project has been very extensive and is included in an 11-experiment paper I sent in a few months ago that is to be published in *Cognitive Psychology*. Here, I will give only a short summary and the direction this project has taken. The short summary is that in many situations people essentially ignore the probabilistic nature of the categorization and act as if it were certain. That is, the likelihood of the alternative categories has little influence on the judgment, under a variety of situations that were used to try to find such an influence. We think this lack of influence derives from a heuristic organization of prediction brought about by basic information

processing limitations (i.e., considering alternatives is too difficult in most situations). We thought that the use of alternative categories in prediction might occur with more naturalistic situations (Malt, Ross, & Murphy, under review at JEP:LMC), but we still found little influence. Our current work (Ross & Murphy, in progress) addresses more extensively the situations under which one might be able to find such an influence. If the alternatives are brought to mind at the time of prediction, either explicitly or by the nature of the prediction judgment, then some influence is found. This project has been important for my problem solving work as well, both in forcing me to think more concretely about how various categories may be used and also in considering the limits of information processing more directly in the learning process.

General Summary

I have presented a brief summary of the results from the projects addressing the three objectives. As is evident, my work continues to address each of these objectives. I have tried to outline for each where I see the work going in the next year. The projects that seem most ripe for advances are the problem categorization work (1ai), the work on category use (1aii), the interaction among processes (1b), and the predictions from categories (3b). The work on problem solving expertise (2) has gone well, but is being finished up, as is the work on reminding-based learning in categorization (3a).

References

Blessing, S. B., & Ross, B. H. (under review). Content effects on problem categorization and problem solving.

Bradshaw, G. L., & Ross, B. H. (1993). Reminding and interpretation during encoding. Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, 249-253, Boulder, CO.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, 5, 121-152.

Hinsley, D. A., Hayes, J. R., & Simon, H. A. (1977). From words to equations: Meaning and representation in algebra word problems. In M. A. Just & P. A. Carpenter (Eds.), Cognitive processes in comprehension. Hillsdale, NJ: Erlbaum.

Malt, B. C., Ross, B. H., & Murphy, G. L. (under review). Predicting features for members of natural categories when categorization is uncertain.

Murphy, G. L., & Ross, B. H. (in press). Predictions from uncertain categorizations. Cognitive Psychology.

Ross, B. H. (1984). The effects of remindings on learning a cognitive skill. Cognitive Psychology, 16, 371-416.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13, 629-639.

Ross, B. H. (1989a). Remindings in learning and instruction. In S. Vosniadou & A. Ortony (Eds.) Similarity and analogical reasoning. Cambridge: Cambridge University Press.

Ross, B. H. (1989b). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 456-468.

Ross, B. H., & Bradshaw, G. L. (in press). Encoding effects of remindings. Memory & Cognition.

Ross, B. H. & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition, 16, 42-55.

Ross, B. H., Perkins, S. & Tenpenny, P. L. (1990) Reminding-based category learning. Cognitive Psychology, 22.

Ross, B. H. & Spalding, T. L. (1991). Some influences of instance comparisons in concept formation. In D. Fisher, M. Pazzani,& P. Langley (Eds.) Concept formation: Knowledge and experience in unsupervised learning. San Mateo, CA: Morgan Kaufmann.

Ross, B. H., & Spalding, T. L. (in press) Concepts and categories. In R. Sternberg, (Ed.), Handbook of Perception and Cognition, Vol. 12, Thinking and Problem Solving. San Diego, CA: Academic Press, Inc.

Spalding, T. L., & Ross, B. H. (in press). Comparison-based learning: Effects of comparing instances during category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition.

PUBLICATIONS ACKNOWLEDGING AFOSR SUPPORT

Ross, B. H., Perkins, S. & Tenpenny, P. L. (1990) Reminding-based category learning. Cognitive Psychology, 22.

Ross, B.H. (1990). The access and use of relevant knowledge: A specific case and general issues. In R. Freedle (Ed.) Artificial intelligence and the future of testing. Hillsdale, NJ: Erlbaum.

Ross, B. H. & Spalding, T. L. (1991). Some influences of instance comparisons in concept formation. In D. Fisher, M. Pazzani,& P. Langley (Eds.) Concept formation: Knowledge and experience in unsupervised learning. San Mateo, CA: Morgan Kaufmann.

Bradshaw, G. L., & Ross, B. H. (1993). Reminding and interpretation during encoding. Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, 249-253, Boulder, CO.

Murphy, G. L., & Ross, B. H. (in press). Predictions from uncertain categorizations. Cognitive Psychology.

Ross, B. H., & Spalding, T. L. (in press) Concepts a:.d categories. In R. Sternberg, (Ed.), Handbook of Perception and Cognition. Vol. 12. Thinking and Problem Solving. San Diego, CA: Academic Press, Inc.

Ross, B. H., & Bradshaw, G. L. (in press). Encoding effects of remindings. Memory & Cognition.

Spalding, T. L., & Ross, B. H. (in press). Comparison-based learning: Effects of comparing instances during category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition.

Blessing, S. B., & Ross, B. H. (under review). Content effects on problem categorization and problem solving.

Malt, B. C., Ross, B. H., & Murphy, G. L. (under review). Predicting features for members of natural categories when categorization is uncertain.

PRESENTATIONS BASED UPON AFOSR FUNDED WORK

Reminding-based learning, AFOSR Workshop on Cognition, Arlington, VA, November, 1989.

Some influences of instance comparisons in concept formation. Invited speaker, Workshop on Computational Approaches to Concept Formation, Palo Alto, CA, January, 1990.

Reminding-based problem solving and learning. Cognitive Science, Georgia Institute of Technology. Atlanta, GA, April, 1990.

Reminding-based generalizations in categorization. Midwestern Psychological Association, Chicago, May, 1990. (Spalding & Ross, presented by Spalding).

Reminding-based problem solving and learning. McMaster University. Hamilton, Ontario. December 1990.

Focussing effects of reminding-based category learning. Midwestern Psychological Association, Chicago, May 1991. (Spalding & Ross, presented by Spalding).

Conservatism in learning. GUV meeting, Bloomington IN, September, 1991.

Concept formation and the learning of feature correlations. (T. Spalding & B. Ross), MPA, May 1992.

Learning from the use of earlier examples: Evidence about what is learned. (M. Preslar & B. Ross), MPA, 1992.

Categorizing and solving algebra word problems. GUV meeting, Ann Arbor, MI, October 1992.

Instance comparisons during category learning. (T. Spalding & B. Ross), Psychonomics Society, St. Louis, November, 1992. [poster]

Reminding-based learning. University of Michigan, Ann Arbor, MI, December, 1992.

Remindings during learning. University of Chicago. Feb 1993.

Concept formation and the learning of feature correlations. (T. Spalding & B. Ross), MPA, Chicago, May, 1993.

Comparison-based category learning. (T.Spalding & B.Ross). Poster presented at the American Psychological Society Convention, Chicago, IL. June, 1993.

Encoding and remindings. GUV meeting, Ann Arbor, MI, October 1993.

The effects of remindings on understanding. University of Michigan, Ann Arbor, MI, October 93.